Teaching Network Architecture through Case Studies

Dan Massey
Colorado State University
massey@cs.colostate.edu

Christos Papadopoulos
Colorado State University
christos@cs.colostate.edu

Lan Wang
University of Memphis
lanwang@memphis.edu

Beichuan Zhang
University of Arizona
bzhang@arizona.edu

Lixia Zhang
UCLA
lixia@cs.ucla.edu

ABSTRACT

Today's Internet architecture is facing significant challenges posed by an ever increasing array of devices and applications, tremendous amount of traffic, as well as serious security threats. There has been considerable debate in the networking community on how to evolve or redesign the Internet. However, our graduate networking courses often favor the mechanical analysis of specific protocol details over understanding of high-level architectural ideas. As a result, our students master the *how*, often at the expense of the *why*. This paper presents our attempt to add more examination of architectural designs into graduate courses on networking. Our premise is that a network architecture is not borne out of a singular piece of work, but rather a progression of ideas that can be traced through a series of papers. We believe it is important for our students to *understand* the progression of these ideas and the design tradeoffs being made, rather than simply *knowing* how the architecture works. To illustrate our approach while limiting new material, we focus on two network architectures, the current Internet and one of the new architectures based on content-centric networking. For each architecture we provide a series of five papers that aim to show both the big ideas and the evolution of their architectural designs.

1. INTRODUCTION

Today's students in the networking area live in an exciting as well as challenging time. At the lower layers, advances in transmission technologies such as optics and wireless offer dramatic changes in the way devices connect to the network. Along with changes in communication media, the devices themselves are also rapidly advancing. The days of connecting large computer centers have long past and today smart phones and sensor nodes are becoming ubiquitous. With proliferation of smart phones, communication devices are not just always on, but they are always on you. At the higher layers, the traditional file transfer applications have expanded to include P2P sharing, streaming services, online gaming, and a vast set of other applications. In the middle, the IP layer itself has been undergoing major changes with the deployment of carrier grade NAT, IP multicast, IPv6, and other proposed additions to the IP layer. Security is no longer a side topic; now any well thought out protocol design must include security considerations. And of course one cannot leave out changes in mobility, scaling, and QoS. In addition, today's networks have direct impact on social issues and no discussion of networking is complete without economic considerations.

There have been considerable debates on how all these changes impact the current network architecture. Some argue these changes should inspire a revolution in network architecture. Just as the Internet completely revolutionized communications networks, the next revolution should usher in a new, perhaps yet to be envisioned, network architecture. Others argue that what is needed is evolution, not revolution: the future network architecture may indeed look quite different from today, but the change will occur as a number of evolutionary steps (and mis-steps). Still others argue that we already have the correct basic network architecture and one just needs to make updates and incremental changes. Regardless of one's position on where the Internet architecture is headed, two facts seem indisputable. First, *today's students will be users of the future Internet and some of them will become network architects of the future*. Second, we can only expect a better network architecture *if students are taught to analyze the design constraints and understand the trade-offs of different designs.*

To motivate and guide students into learning architectural designs, this paper explores the use of two different architecture designs as case studies, and articulates what would be effective ways to incorporate these architectures into a graduate networking course. The architectures we choose to cover are the current Internet architecture and a proposed design of future architecture based on content centric networking. The choice of the Internet architecture is obvious, but not sufficient. We argue that one can teach architecture more effectively by examining multiple architectures. For a second architecture, we selected Named Data Networking since we are actively working on its development. Our choice of a contrasting architecture is based on our experience, and we in no way mean to imply that this is the *right* or even the *only* choice. We hope and believe that the networking community will benefit greatly if other teams working on different architectures develop similar modules. Educators can then select which modules to use based on time, interest and other criteria of their choice.

The remainder of the paper is organized as follows. Section 2 discusses what we hope to achieve in teaching architectures. Section 3 presents the ten papers we think best teach architectural concepts. Section 4 discusses issues we have encountered in covering the papers in a graduate course. Finally, Section 5 concludes the paper.

2. OBJECTIVES OF TEACHING ARCHITECTURES

Our objectives of teaching architectures are two-fold. On one hand, we want to encourage new and novel thinking by students. We are particularly concerned by comments that today's network was somehow a pre-ordained system. It is of course essential to understand the workings of TCP or BGP or DNS or any of the common protocols, but this can lead to a mistaken impression that today's specific protocols and more broadly today's architecture represents a set of fixed points that cannot be replaced by dramatically new designs. On the other hand, we want to convey that designing a new architecture is not a trivial step and there are key lessons that we hope the students will take away from any course.

New architectures do not simply fall from the sky. If we convey nothing else in our courses, we hope to teach that (good) new architectures are not invented in one atomic step. One wants to encourage new and novel ideas, but at the same time one needs to teach that good architectures are developed over time as a series of (sometimes revolutionary) ideas and refinements. One cannot understand an architecture by reading a single"seminal paper" and one does not invent a new architecture in one attempt. To convey this, we must demonstrate architectural development over a series of papers that often span decades.

Architecture involves trade-offs. Developing an architecture requires understanding the objectives and constraints as well as making appropriate trade-offs. At a high level, everyone supports the idea of making the network efficient, secure, robust to failures, easy to manage, and so forth. It is easy to say that the Internet overlooked security at many key points in the design. It is much more challenging to think about what might be lost by prioritizing security over say robustness or automatic configuration. The challenge is that one needs to prioritize goals and further note that some goals may be mutually exclusive.

In order to teach students about network architectures, one needs to discuss *more than one architecture.* Although a particular architecture such as the Internet could by itself be the subject of a course, deeper understanding of architectural principles comes from comparing and contrasting more than one design. To a large extent, we do this today by comparing the Internet architecture with the telephone network. In summary, our challenge is to cover the evolution of an architecture, do this for multiple architectures, and during a limited portion of a semester.

2.1 Proposed Teaching Strategy

The evolution of the Internet architecture provides a rich set of material. While a discussion of the Internet architecture could be a course by itself, our objective here is to provide a module on network architecture, not to provide a complete history of the Internet or identify all key papers leading to its development. Using the Internet as an example, some of our key points become immediately clear. The Internet architecture did not simply fall from the sky and there is no single paper that describes the Internet. Instead one can see the Internet develop over a series of seminal papers. The Internet also faced clear trade-offs and the discussion of these trade-offs help to see both the advantages and disadvantages of today's Internet.

We use the Named Data Networking (NDN) [10] architecture as a second example. Note that the choice of NDN is somewhat arbitrary – we chose it because we are involved in this work, but other architectural designs could be inserted in its place. We do, however, advocate teaching *some* other architecture along with the Internet. A single example is not sufficient to teach any idea and using only a single example blurs the distinction between principles that apply to any architecture (e.g. one must identify and prioritize goals) and the architecturally specific design choices (e.g. which goal was the top choice in a particular architecture).

One could argue that two architectures are too few and inclusions of more architectures would be beneficial. We agree with this statement. However, our goal is to fit this material into an already overloaded curriculum. We further argue that each architecture must be presented as a series of papers leading toward the architecture. It is not sufficient to present a single paper as the description of any architecture. It is the progression toward the architecture that is essential to convey. Toward this end, we propose to cover five papers per architecture. One could also argue that five papers is too few, but at the same time, one could also argue that 10 papers are too many to add in a course. We do not claim the choice of two architectures with five papers for each is ideal, but we do claim it is feasible and we are using this approach in our own courses. We simply hope that this is a starting point for further debate.

3. TEN ARCHITECTURAL PAPERS

We begin our introduction to network architecture by first covering the Internet. Studying the Internet is somewhat easier because today's students can relate to it and thus are interested in the evolution of the network. We then make a dramatic shift from a network focused on the *where* to a network focused on the *what*. We emphasize that the Internet focuses on fetching data from a destination IP address while the NDN architecture eliminates the concept of a source entirely and replaces the destination (currently an address) with the name of the desired data.

3.1 Five Papers on The Internet Architecture

Paper 1: "On Distributed Communications Networks" by Paul Baran [1]. We begin our study in the very early days of packet network architectures. Baran seeks to build a robust network that can survive massive losses. An attack disables many of the nodes and Baran is interested in finding the largest surviving connected component. The networks of the time rely on a highly structured system of central nodes so the destruction of only a few central nodes disconnects the entire network. The paper illustrates the importance of redundancy in building a resilient system. With a small amount of redundancy, Baran shows the topology remains connected even when large numbers of nodes are destroyed. To take advantage of this surviving topology, the network would operate on packets rather than fixed circuits and use dynamic routing of packets. These are novel ideas at the time and start us down the road toward packet network architectures[1].

Paper 2: "A Protocol for Packet Network Interconnection" by V.G. Cerf and R.E. Kahn [2]. Having introduced the concept of packet in the previous paper we now build the simple packet idea into early TCP/IP in this classic paper. At the graduate level, students have already been through an undergraduate course and should have seen many of the concepts of this paper. Among its many features, it introduces gateways for interconnection, retransmission windows, and the thin waist idea of putting only the minimal requirements into interconnection layer. Today, students simply expect that there is not a single link layer technology in the Internet. Most laptops have both wired and wireless interfaces and most smart phones have 3G, WiFi, and bluetooth. However, it was not always clear that devices should interoperate across multiple different link layer technologies. The network architecture made a particular choice to allow a “thin" interconnection layer.

Paper 3: "End-to-end Arguments in System Design" by J. Saltzer, D. Reed, and D. Clark [11]. In the first two papers, we presented the design for a redundant packet network with a thin waist. The result is the well known hourglass design that will be discussed further in paper 4. However, first we need to consider where to place all the additional features that a network clearly needs to support. Reliable delivery, security, and many other services are needed if we are to build a truly viable network. Do we build complexity at the lower part of the hour glass (link layer) or upper part (transport/app layer)? At this point, it is time to introduce the end-to-end principle. Adding features to the lower layer typically is not a complete solution for reliability, security, and a host of other problems. Adding these features as a mandatory feature of every link is not always plausible if we support any link layer technology, but even if we could, the claim is that it often harms services that do not want new features. Furthermore, in Paper 1 Baran has challenged us to build a network that will survive despite massive losses. We can now introduce the concept of fate sharing. If we want a redundant dynamic network that might suffer massive disruptions, it seems wise not to place a tremendous amount of state in middle of the network. If all the state is at the end, we can continue to function as long as the topology and routing provide the ends with some path to reach each other.

Paper 4 "The Design Philosophy of the DARPA Internet Protocols" by D. Clark [3]. This paper first introduces an obvious but often overlooked fact that in any good design one must identify the design objectives and prioritize the objectives. Packets, redundant connections, gateways, and retransmission windows are all great tools. The designers of the Internet did not simply pick some concepts, throw them together, and hope they produced a useful result. In any good design, the designer set out to achieve something and then find the best tools to achieve their objectives. This lesson is critical for any design.

So what was the early Internet design trying to achive? It turns out we are in luck, we have covered the top three on the list: 1) function despite loss of networks/gateways (Paper 1 by Baran); 2) support multiple types of services (Paper 3 on End to End); and 3) accommodate a variety of networks (Paper 2 by Cerf/Kahn). The Internet is notoriously bad at accounting for each individual packet. It also turns out that cost accounting was on the list, but was last. If the design had set out to build a network whose first goal was cost accounting, we would likely have a very different Internet. It is also important to note these goals are not set in stone. It may be the case that a different set of objectives and/or a different ordering of the goals is now in order. This is a great opportunity to challenge students to think of their ordered list of objectives. The key idea is that you need to identify what you want to achieve and prioritize the list (because you likely will not get everything on your list).

Paper 5: "Watching the Waist of the Protocol Hourglass" by Steve Deering [4]. This is a presentation, not a paper. However, it is an excellent way to wrap-up all the concepts we have seen so far and show that the network design is continuing to this day. In Paper 2 by Cerf/Kahn, we built a thin waist at the IP layer. This was a key aspect to achieving our design goals as stated in Paper 4 by Clark. The end-to-end principle (Paper 3) helps guide modifications to the design, suggesting modifications are best suited to the ends, rather than bloating the middle with more and more features and expanding the thin waist.

Deering provides an insightful and entertaining view of temptations that threaten to expand the thin waist. Multicast and QoS are just a few of the features. The presentation also introduces the role of IPv6 and provides insights on why

[1] As a side note, we believe this paper helps explain the occasional statement that “the Internet was designed to survive a nuclear war”. To the best of our knowledge, this was never a stated objective of either the early ARPANET or later Internet. However, Baran’s motivation is to survive a nuclear strike and much of the work on packet networking traces its origins to this paper by Baran.

IPv6 deployment remains such a challenging step for the Internet. Overall, we now have an excellent overview of key architectural ideas and design choices in the current Internet.

3.2 Five Papers on Named-Data Networking Architecture

The existing Internet architecture focuses on communication between source and destination IP addresses. However as the global user community continues to push the frontier of Internet usage, in particular as the number of computing devices increases by leaps and bounds and most of those devices lose their fixed “where" (address) and become mobile, IP’s point-to-point communication model has become a constraint. Now we challenge the students to *delete the concept of location.* Instead, consider that communication is about seeking a specific data item, such as the slides for today’s lecture. The server hosting the data is irrelevant, it does not matter if the slides are provided by the university webserver, the instructor’s personal page, or some other student in the class. It is the data content the consumer seeks that matters, not the location. Of course it does matter that the data is correct. If one receives the slides from another student in the class, one wants to ensure they are the complete and correct slides. Thus this design can only make sense if one has security as part of the design. This provides a starting motivation for Named-Data Networking Architectures.

Paper 6: "Multicast routing in internetworks and extended LANs" by Deering and Cheriton [5]. This paper proposed IP multicast, a brand new communication model. This model greatly facilitates distributed applications that must send the same data to multiple destinations, or those that must locate or query content when the exact location of that content is unknown. As the last hop of communications becomes increasingly wireless to aid mobility, this model can also make best use of the broadcast nature of wireless media.

Paper 7: "A Reliable Multicast Framework for Light-weight Sessions and Application Level Framing" by Floyd, Jacobson, McCanne, Liu, and Zhang [6]. This paper promoted the receiver-based model of data reliability and application level naming.

Paper 8: "Adaptive web caching: towards a new global caching architecture" by Michel, Nguyen, Rosenstein, Zhang, Floyd, and Jacobson [9]. With the addition of web caching, one begins to see a movement away from server locations and toward designs focused on the data itself. This early work is not yet introducing a fundamentally new design, but is introducing the concept of designing a network whose primary service is access to data rather than access to a particular server.

Paper 9: "Building Efficient Wireless Sensor Networks with Low-Level Naming" by Heidemann, Silva, Intanagonwiwat,, Govindan, Estrin, and Ganesan [7]. New architectural designs need driving motivations. Even an ideal design will not succeed in practice without external factors that motivate the deployment of the new architecture. In the case of the Internet design, the architecture evolved along with the introduction of the PC and the move from massive computer centers to desktops and then laptops. To support sensor networking this paper moved the notion of application framing to a new stage, and designed an efficient delivery system that is based on *named data* and in-network processing.

Had the Internet designed in an age of sensor networks and smart devices, would one still develop the same network designs and naming conventions? The concept of location specific naming and computation is challenged by the large numbers of wireless devices. When combined, papers 8 and 9 introduce new challenges and new directions for both familiar web content and new sensor network inspired designs. A common theme in both papers is a challenge to the traditional location based network design. In both paper 8 and 9, the idea of providing communication to a specific IP address is not well matched with the overall objectives of the data consumers.

Paper 10: Networking Named Content, by Jacobson, Smetters, Thornton, Plass, Briggs, and Braynard [8]. We conclude our list by what we feel is an exciting new direction. Even if one disagrees entirely with this direction, it does serve the purpose of forcing students to think differently. In particular, the paper observes that communication in the previous work has focused on a source communicating with a destination. In the current designs, you download a file or fetch a web page from a particular server. Location is an essential part of the discussion, but does it need to be? Jacobson et al. challenge the reader to think in terms of *what and not where*. The location becomes irrelevant. Instead one seeks a particular content. Perhaps the content is located on the data publisher’s website, but it may also be cached anywhere throughout the network. The location is irrelevant.

In this new design, the basic IP packet is replaced by Interest packets and Data packets. These packets have *no source address and no destination address*. This challenges the students to think in a new way. Can a network even operate if packets do not specify a source and destination? What lessons from the above papers change? What stay the same? Paper 3 by Clark still teaches us we need a priority list. Furthermore, his top three goals still stay at or near the top of the list in the new NDN design. We still need thin waist (Paper 2 by Cerf/Kahn and Paper 5 by Deering), but now the waist identifies data name rather than location name. Baran’s (Paper 1) lesson on redundancy is as critical as ever, if not even more important now. The End-to-End Principle (Paper 3) still plays a key role in deciding where the complexity lies, but implementation changes a little from the “end location" to the data publisher and even data itself.

4. CURRICULUM CONSIDERATIONS

In our experience, there are two key pitfalls that have been raised by students and need to be addresses for a successful course.

The first challenge is that some students have viewed the initial papers as *old and dated*. Some students come into the course excited by the latest trend in say 4G wireless networks. They are then stunned to start with a Baran's paper from the 1960s and further appalled by evaluation numbers in the paper that talk in terms of *bits per second*. We have found two items are essential to address these concerns. First, it is important for the lecturer to convey that we are teaching ideas, not history. The design ideas from the early Internet are fundamental to understanding the design of today's newest trend. Second, it is important to provide the student with the broader course schedule so the students are aware early on that the course may begin in the 1960s, but will also look at today's trend.

The second challenge is that these papers can also be viewed as philosophical instead of technical. As a community, we have tended to do an excellent job on teaching students to use and learn packet formats. Students are familiar with the use of the SYN bit in a TCP header, the fragmentation fields in IP packet, and query ID in DNS packets. These are valuable and useful things to know, but learning the packet format from Baran is clearly not valuable. The idea here is to teach architectural concepts and the evolution of design. One needs to explicitly remind the students the exact format of the packet is not the objective. Short quizzes can help further emphasize key design questions and downplay detailed specifics.

On a more positive note, the lecturer needs to stress the important ideas and thinking big. Do you agree with the order in Clark's Design Philosophy? How different would today's network be if say "Cost Accounting" was ranked above "Robustness"? In Jacobson's content centric design, what are the security implications of having a data name but no notion of data location? The course has tended to fail if students believe they are being asked to understand now obsolete packet formats, but succeed in cases where students are challenged to think big.

5. CONCLUSIONS

Today's graduate students will be tomorrow's network architects. Regardless of whether one believes the future Internet will be revolution or an evolution of the current design, it is essential for graduate students to consider architectural issues and have some understanding of how different architectures develop.

Toward this end, this paper proposes teaching network architecture by considering the development of two different designs, the current Internet and a new data centric approach to networking. For each design, we introduce a sequence of five papers that show keys aspects and the evolution of architectural thinking. All the papers here provide big ideas, not all of which you may agree with. Our objective is not to declare these to be the two best architectures, but rather to promote the discussion of architectural design. We hope this leads to more discussion on what other papers should be on the list for thinking architecturally.

6. REFERENCES

[1] P. Baran. On distributed communications networks. *IEEE Transactions on Communications*, March 1964.

[2] V. G. Cerf and R. E. Khan. A protocol for packet network intercommunication. *IEEE TRANSACTIONS ON COMMUNICATIONS*, 22:637–648, 1974.

[3] D. Clark. The design philosophy of the DARPA Internet protocols. In *SIGCOMM*, 1988.

[4] S. Deering. Watching the Waist of the Protocol Hourglass, 2001. IETF51 Plenary Talk.

[5] S. E. Deering and D. R. Cheriton. Multicast routing in datagram internetworks and extended lans. *ACM Transactions on Computer Systems*, 8:85–110, 1990.

[6] S. Floyd, V. Jacobson, C. gung Liu, S. Mccanne, and L. Zhang. A reliable multicast framework for light-weight sessions and application level framing. In *IEEE/ACM Transactions on Networking*, pages 342–356, 1995.

[7] J. Heidemann, F. Silva, C. Intanagonwiwat, R. Govindan, D. Estrin, and D. Ganesan. Building efficient wireless sensor networks with low-level naming. In *Proceedings of the eighteenth ACM symposium on Operating systems principles*, SOSP '01, pages 146–159. ACM, 2001.

[8] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard. Networking named content. In *Proceedings of the 5th international conference on Emerging networking experiments and technologies*, CoNEXT '09, pages 1–12, New York, NY, USA, 2009. ACM.

[9] S. Michel, K. Nguyen, A. Rosenstein, L. Zhang, S. Floyd, and V. Jacobson. Adaptive web caching: towards a new global caching architecture. *Comput. Netw. ISDN Syst.*, 30:2169–2177, November 1998.

[10] Named Data Networking. `http://www.named-data.org`.

[11] J. H. Saltzer, D. P. Reed, and D. D. Clark. End-to-end arguments in system design. *ACM Transactions on Computer Systems*, 2:277–288, 1984.